Patterns of Connection

Matthew Blumberg

Background

My interest in Human Computation– described here as "Distributed Thinking"– dates back to 2008 and the FIFA World Cup final. The truth is (being American) I didn't watch. I only read about it the next day. It was quite a match, apparently– eventually won 1-0, on a 73rd-minute goal by Wayne Rooney of Manchester United. But what was most notable in the coverage– to me– was the comment that the match had been watched, live, by 700 million people.

A soccer game being about 90 min long, this amounts to more than *a billion hours* of human attention– focused on a bouncing ball. That's about 120,000 *person-years* of attention– compressed into 90 min.

Which raised the question: what could be done with all that cognition? Could it be harnessed for constructive purposes? What knowledge and tools and methods would be required?

Crowdsourcing

A number of web-based projects have emerged which draw on the aggregated intellectual skills of large numbers of people over the Internet. These projects represent the "state of the art" in Human Computation—exciting efforts to harness many minds in order to do intellectual work that would otherwise be impossible. A few key examples follow (there are of course many others):

M. Blumberg (✉)
GridRepublic, USA
e-mail: www.gridrepublic.org

P. Michelucci (ed.), *Handbook of Human Computation*,
DOI 10.1007/978-1-4614-8806-4_2,

- *Clickworkers (2001)*—People were shown images of the surface of Mars, and asked to help map it by drawing circles around the craters. (Computers aren't good at this sort of pattern recognition, but people are.[1])
- *Stardust@home (2006)*—A NASA probe dragged a volume of gel through the tail of a comet; the comet particles were quite few and small, and searching for them in the large volume of gel was a challenge. The Stardust team posted nearly a million images of small sections of the volume online, and people were asked to search through these and to find characteristic tracks of particles. This collective effort considerably accelerated the search for the "needles" in the "haystack"
- *Galaxy Zoo (2007)*—People are shown images of galaxies, and asked to categorize them by visual features: spiral, disk, etc.; the goal is to build a celestial almanac. (As above, computers aren't good at this sort of image analysis.)
- *ESP Game (2003)*—Pairs of people are shown an image at the same time, and each starts typing descriptive words. When both have entered the same word, they "win" (and the system presumes to have learned a useful "tag" for use in categorizing the image).
- *Ushahidi (2008)*—People in and around crisis situations submit reports by web and mobile phones. These are aggregated (and organized temporally and geospatially), to give an accurate and unmediated view of the emerging situation.[2]
- *eBird (2002)*—Bird watchers throughout the world submit observations, creating a real-time database of bird distribution and abundance.
- *Iowa Electronic Market (1995)*—People buy and sell "contracts" in a (not-for-profit) Futures market, as a tool for predicting outcomes of elections, Hollywood box office returns, and other cultural phenomena.
- *FoldIt (2008)*—People solve 3D visual puzzles, as a means to solve problems in protein structure prediction.
- *Phylo (2010)*—People search for matching patterns in sequences of DNA, represented as strings of colored blocks.
- *EteRNA (2010)*—People solve visual puzzles related to the folding of RNA molecules

The above represents a fairly wide range of objectives and activities—thought it may be observed that all follow a certain pattern, one which is presently characteristic of what is commonly referred to as Crowdsourcing:

- In each project above, all users perform the same task repetitively (i.e., all users draw circles to mark craters, or place a pin to mark traces of comet, or find matching patterns in strings of colored blocks.)
- In most cases, the task is quite simple; it is the vast quantity that must be slogged through which requires the crowd input.

[1] Possibly of interest: see article in this volume by Jordan Crouser and Remco Change, discussing relative strengths of humans vs computers.

[2] Possibly of interest: See article in this volume on crowdsourcing disaster relief by Ushahidi founder Patrick Meier. Human Computation for Disaster Response.

- Tasks are single-user: interaction among participants while performing the work is not required.[3]
- There is no parceling of task-type based on user expertise (At most, users of measured skill—ie users who have returned validated results—might get harder versions of the task at hand.)

In sum, with these tasks, there is no "higher level" thinking being done by the "Crowdsourcing" system. All of the tasks completed by the public (individually and collectively) could plausibly have been done by the project organizers—in most cases better.[4] The projects are really a means of collecting and applying large quantities of unskilled labor. This of course is useful; but much more is possible.

The discussion below seeks to make the case that it is possible to create "Thinking" systems—systems created of many minds, and capable of sophisticated problem solving….

Distributed Thinking

In order to contemplate what a large scale thinking system might look like, it is useful to have a notion of what *Thinking* is.

As a point of reference, consider the model proposed by Marvin Minsky in *Society of Mind* (1988). In Minsky's model "minds are built from mindless stuff".

Minsky hypothesizes that a Mind—that thinking—is made up of many small processes (which he calls "agents"); that these are simple; that they are not especially intelligent in and of themselves—And that *it is the way that these things are connected* that creates intelligence, as a sort of emergent property of the "thinking" system.

Picking Up a Cup of Tea

For example, if one wanted to pick up a cup of tea there might be several processes involved (several "agents"):

- Your GRASPING agents want to keep hold of the cup
- Your BALANCING agents want to keep the tea from spilling
- Your THIRST agents want you to drink the tea
- Your MOVING agents want to get the cup to your lips

[3] ESP game is an exception here; sort of.

[4] A notable exception is FoldIt: In the case of FoldIt, it turned out that a public participant was unusually good at the task, better than subject area experts. This fact alone highlights the sophistication of that project. I.e., FoldIt serves to demonstrate the example that when projects are sufficiently advanced, they may draw in "savants", persons unusually good at the particular task—better in some cases than the project organizers themselves. And/or, projects may empower novel combinations of intellectual skills of persons otherwise unknown the project organizers.

… These would all be independent processes, performed in parallel, competing for resources in various ways—and collectively producing the behavior of picking up and drinking the cup of tea.

Stacking Blocks

Another illustration, a slightly more complicated cognitive problem– Imagine you had a pile of blocks, and you wanted to pile them up in a stack. You might hypothesize the existence of a "mental program" to do this, call it "Builder" (Fig. 1):

Fig. 1

In the Minsky view of the mind, this program would be composed of smaller applications, for instance (Fig. 2):

Fig. 2

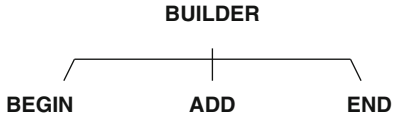

And each of these "programs" or "agents" would themselves be composed of smaller functions. And each of these, of possibly smaller… Until you got down to some list basic "primitive" functions from which all the others are built (Fig. 3):

Fig. 3

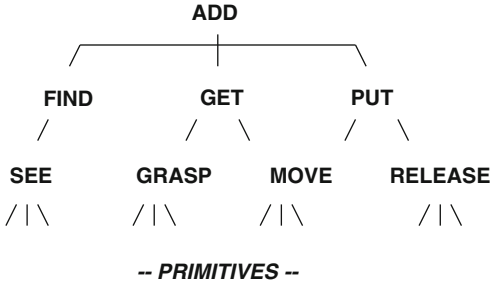

What's interesting about this approach is that if you took from the previous chart describing "Builder" only the list of the Agents themselves, you wouldn't know anything about what the Builder does. It's only when you put the things into a structure that it becomes possible to contemplate that they might do something useful (Fig. 4):

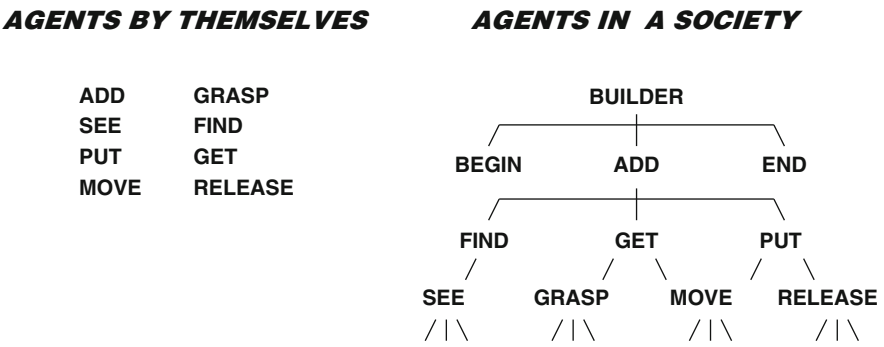

Fig. 4

This brings us to the first essential point of this essay: *Intelligence is created not from intellectual skill, but from the patterns within which intellectual skills are connected.*

The Minsky "Society of Mind" model is but one example; in general, patterns of organization which result in emergent "intelligent" behavior may be referred to as *"Cognitive Architectures"*.

From Crowdsourcing to Intelligent Systems

With an eye towards imagining a system which has a higher level of intelligence than its individual participants, and following Minsky's Cognitive Architecture– it's perhaps interesting to imagine what the set of "primitives" (the basic, unintelligent functions from which more complicated processes might be built) could be. Perhaps:

- **Pattern Matching/Difference Identification**
- **Categorizing/Tagging/Naming**
- **Sorting**
- **Remembering**
- **Observing**
- **Questioning**
- **Simulating/Predicting**
- **Optimizing**
- **Making Analogies**
- **Acquiring New Processes**

…This is not meant as a comprehensive list, just some illustrative examples. Note that none of these functions are especially complicated in and of themselves (though several are to varying degrees computationally intractable). Most are, in a wide range of contexts, quite parallelizable.

As food for thought, consider that many of the previously listed crowdsourcing projects provide quite nice templates for several of these very activities:

- **Pattern Matching/Difference Identification—**As noted, in *Clickworkers*, participants identified circles in a database of images; in *Stardust@home*, participants identified characteristic traces of comet dust in a database of images; in a range of other projects participants mark features on satellite images to generate or enrich maps, etc.
- **Categorizing**—In *Galaxy Zoo*, participants are shown images of galaxies, and asked to categorize them, by visual features: spiral, disk, etc.—and this is used to build up a structured database of astronomical objects.
- **Tagging/Naming—**In *ESP Game* participants create useful tags for image search (*In fact the system was licensed by Google to improve their image-search functionality).
- **Observing—**In *Ushahidi*, in *eBird*, and many other projects, distributed observations are entered into a shared central database
- **Simulating/Predicting—**In *Iowa Electronic Market*, and a wide range of subsequent "Prediction Markets", participants engage in a process which has been shown to effectively predict the outcome of a range of events.
- **Optimizing—**In *FoldIt* participants are asked to optimize the shape of an object according to certain parameters.
- **Etc…**

Following the earlier discussion, while it may be the case that any individual one of these systems is useful and interesting, it is the potential of *putting these things together into systems*– into intelligent patterns, into Cognitive Architectures– where really interesting things may become possible.

A Speculative Example

Imagine creating a drug discovery pipeline using Distributed Thinking –

By way of context, note that one method of drug discovery is {1} to identify a mutant or malformed protein which has been implicated in a specific pathology. And then {2} to find some other protein that binds to this deviant but nothing else– this is akin to sticking a monkey wrench into a running machine: the goal is to muck up the works, to cause that process to fail. And this can be quite effective.

Given a target identified by lab work, one could imagine subsequently breaking the process of discovering such "monkey-wrench" proteins into a sequence of steps– like, "docking" to see what candidate proteins stick to your target; "similarity analysis" to see which proteins are like which other proteins (to find alternative avenues of exploration); "optimizing" (to improve marginally useful candidates); "cross screening" (to see if a candidate has side effects, by checking whether it docks with anything it's not supposed to); and so on… (Fig. 5).

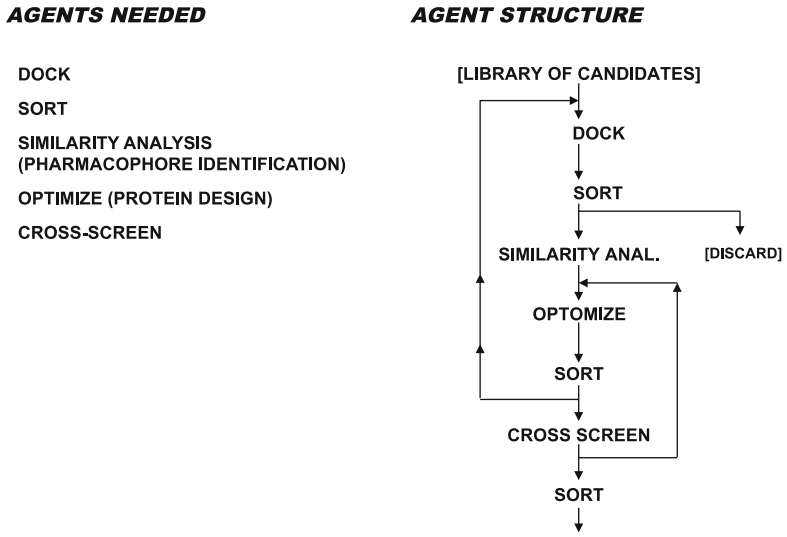

Fig. 5

All of these individual steps/processes could be imagined in terms of systems on par with existing crowdsourcing applications. And one could imagine linking these functions—these agents– into a fairly elaborate workflow,[5] the collective function of which would be to seek out and create promising drug candidates.[6]

Summary

A great deal has been done with Crowdsourcing. Current examples share a number of features however, most notably insofar as each supports only a single type of task. The discussion presents the idea of "Cognitive Architectures"—patterns into which individual systems, each performing specific (and potentially mundane) tasks, might be interconnected to collectively create a higher level of cognition.

[5] A small but important step in the evolution from current-generation Crowdsourcing to Distributed Thinking would be adoption of a standard means to integrate individual projects (individual functions) into more complex workflows. It is hoped that developers of such projects—and especially developers of middleware like BOSSA and PyBOSSA—will provide APIs that enable others to submit inputs and collect outputs, so the output of one project might be used as the input for another. For instance, a Phylo-style DNA project might input sequences into a FoldIt style structure prediction application.

[6] Of course there are numerous technical reasons why these steps are incomplete or may be impractical or infeasible at the moment. As noted it is a speculative example, a broad-stroke illustration.

The goal is to raise the prospect that "Apollo Project" challenges might be met by the application of sufficient attention, properly structured—It's all a matter of the patterns by which we connect ourselves and our information.

References

Minsky M (1988) The society of mind. Simon and Schuster, New York. ISBN 0-671-65713-5

Websites

Clickworkers: Originally http://clickworkers.arc.nasa.gov/top. Later revised to http://clickworkers.arc.nasa.gov/hirise. Current version: http://beamartian.jpl.nasa.gov/welcome
Stardust@home: http://stardustathome.ssl.berkeley.edu
Galaxy Zoo: www.galaxyzoo.org
ESP Game: originally: http://www.espgame.org; subsequently included in images.google.com
Ushahidi: www.ushahidi.com
eBird: ebird.org
Iowa Electronic Markets—tippie.uiowa.edu/iem
FoldIt—fold.it
Phylo—phylo.cs.mcgill.ca/eng
EteRNA—eterna.cmu.edu/

http://www.springer.com/978-1-4614-8805-7

Handbook of Human Computation
Michelucci, P. (Ed.)
2013, XLI, 1059 p. 251 illus., 209 illus. in color.,
Hardcover
ISBN: 978-1-4614-8805-7